

A Lower Bound for Light Spanners in General Graphs*

Greg Bodwin and Jeremy Flics
University of Michigan EECS
`{bodwin,jflics}@umich.edu`

Abstract

A recent upper bound by Le and Solomon [STOC '23] has established that every n -node graph has a $(1 + \varepsilon)(2k - 1)$ -spanner with lightness $O(\varepsilon^{-1}n^{1/k})$. This bound is optimal up to its dependence on ε ; the remaining open problem is whether this dependence can be improved or perhaps even removed entirely.

We show that the ε -dependence cannot in fact be completely removed. For constant k and for $\varepsilon := \Theta(n^{-\frac{1}{2k-1}})$, we show a lower bound on lightness of

$$\Omega\left(\varepsilon^{-1/k} n^{1/k}\right).$$

For example, this implies that there are graphs for which any 3-spanner has lightness $\Omega(n^{2/3})$, improving on the previous lower bound of $\Omega(n^{1/2})$.

An unusual feature of our lower bound is that it is conditional on the girth conjecture with parameter $k - 1$ rather than k . We additionally show that this implies certain technical limitations to improving our lower bound further. In particular, under the same conditional, generalizing our lower bound to all ε or obtaining an optimal ε -dependence are as hard as proving the girth conjecture for all constant k .

1 Introduction

We study *spanners*, which are fundamental graph sparsifiers with applications in networking, chip design, flow algorithms, etc. [1]

DEFINITION 1. (SPANNERS [13, 14]) *Given an input graph G , a k -spanner is a subgraph H satisfying $\text{dist}_H(s, t) \leq k \cdot \text{dist}_G(s, t)$ for all nodes s, t .*

The goal is generally to construct spanners with a favorable tradeoff between their stretch k and their size. There are two common ways to quantify the size of a spanner H . The first is by its *sparsity*, that is, the goal is to minimize the number of edges $|E(H)|$. The general stretch/sparsity tradeoff is well understood, thanks to a classic result of Althöfer et al. [3]:

THEOREM 1.1. ([3]) *For all positive integers n, k , every n -node graph has a $(2k - 1)$ -spanner on $O(n^{1+1/k})$ edges.*

There is also a matching lower bound, conditional on the *girth conjecture* of Erdős. With respect to a parameter k , which may be any positive integer, that conjecture is:

CONJECTURE 1. (GIRTH CONJECTURE [10]) *There exists a family of n -node graphs with $\Omega(n^{1+1/k})$ edges and girth $> 2k$. (The girth of a graph is the least number of edges in any of its cycles.)*

If one removes any edge (u, v) from a girth conjecture graph, there can be no alternate $u \rightsquigarrow v$ path of length $\leq 2k - 1$ (or else the graph has a short cycle). Thus, any graph G from the girth conjecture has no $(2k - 1)$ -spanner except for G itself, which means that G gives a matching lower bound to Theorem 1.1. The girth conjecture is confirmed for

$$k \in \{1, 2, 3, 5, \Omega(\log n)\},$$

*This work was supported by NSF:AF 2153680.

Stretch	Lightness	Citation
$2k - 1$	$O(n/k)$	[3]
$(1 + \varepsilon) \cdot (2k - 1)$	$O_\varepsilon(k \cdot n^{1/k})$	[7]
$(1 + \varepsilon) \cdot (2k - 1)$	$O_\varepsilon\left(\frac{k}{\log k} \cdot n^{1/k}\right)$	[9]
$(1 + \varepsilon) \cdot (2k - 1)$	$O(\varepsilon^{-(3+2/k)} n^{1/k})$	[8]
$(1 + \varepsilon) \cdot (2k - 1)$	$O(\varepsilon^{-1} n^{1/k})$	[12] (see also [4])
$2k - 1$	$\Omega(n^{1/k})$	conditional on $\text{GC}(k)$ [10]
$(1 + \varepsilon) \cdot (2k - 1)$, for fixed k and $\varepsilon = \Theta(n^{-\frac{1}{2k-1}})$	$\Omega(\varepsilon^{-1/k} n^{1/k})$	conditional on $\text{GC}(k-1)$ (this paper)

Table 1: Summary of prior work on light spanners. In the lower bounds, $\text{GC}(k)$ refers to the girth conjecture with parameter k . See also [11] for discussion of existentially optimal spanner algorithms, and [2, 5, 6] and references within for work on light spanners in various special classes of graphs.

but it is a major open problem to prove or refute it for any other values of k [15, 16].

The other popular way to measure the size of a spanner H is by its *total edge weight* $w(H)$. Since the weight of a spanner can be unbounded, one typically normalizes by the weight of a minimum spanning tree (mst) of the input graph. This measure is called the *lightness* of the spanner.

DEFINITION 2. (SPANNER LIGHTNESS) *The lightness of a spanner H of a graph G is the quantity*

$$\ell(H \mid G) := \frac{w(H)}{w(\text{mst}(G))}.$$

We will also write $\ell(H) := \ell(H \mid H)$ as a shorthand.

Lightness has proved much harder to understand than sparsity; after a long line of inquiry (see Table 1), the community has only recently obtained a stretch/lightness tradeoff that nearly matches the one known for sparsity, up to an additional $(1 + \varepsilon)$ factor in the stretch. The current upper bound is:

THEOREM 1.2. ([4, 12]) *For all positive integers n, k and all $\varepsilon > 0$, every n -node graph G has a $(1 + \varepsilon) \cdot (2k - 1)$ -spanner H of lightness*

$$\ell(H \mid G) \leq O\left(\varepsilon^{-1} \cdot n^{1/k}\right).$$

Meanwhile, on the side of lower bounds for lightness, essentially nothing is known. A girth conjecture graph implies a lower bound of $\Omega(n^{1/k})$ for the lightness of $(2k - 1)$ -spanners, in exactly the same way as for sparsity.¹ But this lower bound does not take advantage of the weighted nature of light spanners, and to date, no lower bounds have been discovered that do leverage edge weights. So it remains conceivable that *no* ε -dependence is really needed, and that the right bounds for spanner sparsity and lightness are the same, in the sense that girth conjecture graphs are the worst case for each.

The main result of the current paper is that this is not the case, and some ε -dependence in lightness bounds is necessary. We show:

THEOREM 1.3. (MAIN RESULT) *Let $k \geq 2$ be a constant positive integer and assume the girth conjecture with parameter $k - 1$. Then there is a family of n -node weighted graphs G for which, letting $\varepsilon := \Theta(n^{-\frac{1}{2k-1}})$, any spanner H of stretch $(1 + \varepsilon)(2k - 1)$ has lightness*

$$\ell(H \mid G) \geq \Omega\left(\varepsilon^{-1/k} n^{1/k}\right).$$

¹By the same argument, an (unweighted) girth conjecture graph G has no nontrivial $(2k - 1)$ -spanner, and its lightness is

$$\ell(G) = \frac{w(G)}{w(\text{mst}(G))} = \frac{|E(G)|}{n - 1} = \Omega(n^{1/k}).$$

Thus, some ε -dependence is needed in general. We can also extend our lower bound to some mildly super-constant k , with some additional technical nuances; we discuss these in Section 2.4. To further help interpret our main result, it implies the following corollary:

COROLLARY 1.1. *For any positive integer n and constant positive integer $k \geq 2$, assuming the girth conjecture with parameter $k - 1$, there exists an n -node weighted graph G for which any $(2k - 1)$ -spanner H has lightness*

$$\ell(H \mid G) \geq \Omega\left(n^{\frac{1}{k} + \frac{1}{k(2k-1)}}\right).$$

The additional factor of $n^{\frac{1}{k(2k-1)}}$ in this corollary is the part that improves over the previous lower bound from the girth conjecture. Considering $k = 2$ for example, this corollary states that there are graphs on which any 3-spanner has lightness $\Omega(n^{2/3})$, improving over the previous lower bound from the girth conjecture of $\Omega(n^{1/2})$.

1.1 Relationship to the Weighted Girth Conjecture The technique used by Althöfer et al. in their proof of Theorem 1.1 is to show that one can always construct a $(2k - 1)$ -spanner of girth $> 2k$, and then to apply the folklore *Moore bound* from extremal graph theory stating that any graph of girth $> 2k$ can have only $O(n^{1+1/k})$ edges. Elkin, Neiman, and Solomon [9] formalized the analogous approach for the light spanner problem by introducing *weighted girth*:

DEFINITION 3. (WEIGHTED GIRTH) *For a cycle C in a weighted graph $G = (V, E, w)$, its normalized weight is the quantity*

$$w^*(C) := \frac{w(C)}{\max_{e \in C} w(e)}.$$

The weighted girth of G is the minimum value of $w^(C)$ over all cycles C in G .*

They then established the following reduction:

THEOREM 1.4. ([9]) *For all $t \geq 1$, every graph has a t -spanner of weighted girth $> t + 1$. Moreover, any graph of weighted girth $> t + 1$ has no nontrivial t -spanner.*

This theorem implies that the question of the lightness needed for a t -spanner is completely equivalent to the question of the maximum possible lightness $\ell(H)$ over all graphs of weighted girth $> t + 1$. We will adopt this reframing of the problem in our lower bound construction and analysis. In particular, to prove Theorem 1.3, our goal will be to construct a graph H of weighted girth $> (1 + \varepsilon) \cdot 2k$ and lightness

$$\ell(H) \geq \Omega\left(\varepsilon^{-1/k} n^{1/k}\right).$$

Going a bit further, Elkin, Neiman, and Solomon conjectured:

CONJECTURE 2. (WEIGHTED GIRTH CONJECTURE [9], C.F. CONJECTURE 1) *For all positive integers n, g , there is an **unweighted** n -node graph that maximizes lightness over all graphs of weighted girth exactly g .*

Our new lower bound does **not** refute this conjecture, and in our view, settling it remains the main open problem in the area of light spanners. However, we do refute certain stronger versions of the conjecture, which may shed light on how best to attack the weighted girth conjecture going forward.

COROLLARY 1.2. (REFUTATION OF “STRONG” WEIGHTED GIRTH CONJECTURES) *The weighted girth conjecture is false if we instead consider the maximum possible lightness over all graphs of weighted girth strictly larger than g , or if we allow g to be non-integral.*

To further explain our point here, let us focus on the setting $g = 4$. It is known that the densest possible graph of unweighted girth exactly 4 has $\Theta(n^2)$ edges (the biclique), but that the densest possible graph of unweighted girth > 4 has $\Theta(n^{3/2})$ edges (the finite projective plane incidence graph [16]). Our lower bound implies that there are graphs of weighted girth > 4 and lightness $\Omega(n^{2/3})$, which beats the lightness bound from the latter construction but not the former. Thus we refute the “strong” weighted girth conjectures quoted above but not the original weighted girth conjecture. A similar effect holds for any even $g \geq 4$.

1.2 Limitations to Further Progress and Technical Overview The ε -dependence in our lower bound is still a far cry from the dependence in the current upper bounds, and so one might hope to improve it further. Specifically, one might hope to generalize our lower bound to all ε , or to improve the dependence on ε to match the dependence in the upper bound. However, due to the fact that our lower bound relies on the girth conjecture with parameter $k - 1$ rather than k , it turns out that either of these improvements will be difficult to achieve in the sense that they would imply the girth conjecture itself.

THEOREM 1.5. (TECHNICAL LIMITATIONS) *Suppose that Theorem 1.3 can be improved in either of the following two ways, while still only conditioning on the girth conjecture with parameter $k - 1$:*

- *It holds for all $0 < \varepsilon < 1$, rather than just $\varepsilon = \Theta(n^{-\frac{1}{2k-1}})$.*
- *It holds with a dependence of ε^{-1} in the lightness lower bound, rather than $\varepsilon^{-1/k}$.*

Then the girth conjecture is true for all constant k .

The girth conjecture has not been settled in any new cases since the '50s [15], and it is regarded as a major open problem in extremal combinatorics and theoretical computer science. So this barrier to improving our approach is likely a serious one; it will probably require major new ideas to prove the girth conjecture for all constant k (if this is even possible).

In light of this barrier, we will next do some technical overviewing of why our lower bound conditions on the girth conjecture with parameter $k - 1$, and the (informal) conceptual challenge in instead conditioning on parameter k . Perhaps the most natural attempt towards a lower bound for light $(1 + \varepsilon)(2k - 1)$ -spanners would be:

1. Begin with a girth conjecture graph G with parameter k ,
2. Increase its edge weights to some value W , and
3. Fix an **mst** with $\Theta(n)$ nodes and edge weight 1, and then carefully embed the nodes of G somehow into the nodes of the **mst**.

The lightness of a graph G obtained from this method would be

$$\ell(G) = \Omega(Wn^{1/k}).$$

Thus we would get a meaningful lower bound so long as W is a favorable function of ε , *and* if we can argue that the final graph G does not admit a nontrivial $(1 + \varepsilon)(2k - 1)$ -spanner. But this seems to be quite difficult. Even though the *initial* graph G has high (weighted or unweighted) girth and thus has no nontrivial spanner, when we embed it into the **mst**, we create new paths between its nodes that might ruin its high weighted girth. So we would need to perform this embedding quite carefully. This would likely require a nuanced structural understanding of G , but since G is conjectured rather than explicit, such a structural understanding seems hard to obtain. Even in the cases where the girth conjecture is confirmed and we do have explicit constructions of G , in particular $k \in \{2, 3, 5\}$, it is still not clear if this embedding can be achieved.

The strategy in this paper is to begin with G as a girth conjecture graph with parameter $k - 1$ instead of k . This has advantages and disadvantages relative to the method sketched above. The main disadvantage is that G will begin with many $2k$ -cycles, which we are now responsible for destroying during our construction. But the main advantage is that G begins with some extra density that we can spend throughout the construction. We essentially spend this by letting our **mst** have $\gg n$ nodes. This allows us to ensure that no *new* problematic cycles are created in the embedding, roughly by including long *spacers* on the **mst** that separate the embedding regions of different nodes of G . It also lets us use a random node-splitting method as we embed to destroy the initial $2k$ -cycles in G . The only new structural understanding of girth conjecture graphs required by this method is a bound on the initial number of $2k$ -cycles in a parameter $k - 1$ girth conjecture graph, which we obtain in Section 2.1.

2 Lower Bound for Light Spanners

In the introduction we fixed k to be a constant and ε to be a particular function of n and k . We will perform most of our construction and analysis with respect to arbitrary $k \leq O(\log n)$ and $0 < \varepsilon < \frac{1}{2}$, only discussing further restrictions of these parameters once they arise. We will assume where convenient that ε^{-1} is integral, which affects our bounds only by lower-order terms.

2.1 The Structure of Girth Conjecture Graphs Before we begin our main construction and analysis, we will need to establish a few structural facts about girth conjecture graphs.

LEMMA 2.1. *For any parameter k , if the girth conjecture holds, then it can be specifically satisfied by a family of graphs G with all of the following properties:*

1. G is bipartite,
2. G is approximately regular (meaning that all nodes have degree $\Theta(d)$ for some parameter d), and
3. For any edge $e \in E(G)$ and any integer $c \geq 1$, the number of $2k + 2c$ cycles in G that contain e is at most $O\left(n^{\frac{k+2c-1}{k}}\right)$.

Proof. Let $G = (V, E)$ be an arbitrary graph satisfying the girth conjecture. Enforcing the first two properties (bipartiteness and approximate regularity) is standard, so we recap the proof somewhat briefly here.

First, to enforce bipartiteness, we take a random partition $V = V_1 \cup V_2$ of the vertex set by placing each node in V_1 or V_2 with equal probability. Then we delete all edges from G that have both endpoints in V_1 or both endpoints in V_2 . Each edge survives with probability $1/2$, and so in expectation we remove only half the edges from E in this step, and we have made G bipartite.

Next, to enforce approximate regularity, fix $d = \Theta(n^{1/k})$ as the current average degree of G . While possible, perform either of the following two steps:

- Find a node of degree $\leq d/4$ and delete it from G , or
- Find a node of degree $\geq d$ and split it into two new nodes, with its incident edges equitably partitioned between the new nodes.

Since we remove at most n nodes in the first step, we remove at most $n \cdot d/4 \leq |E|/2$ edges from G in total, so the total number of edges in G again decreases by at most a constant factor. When the process is complete the average degree in G is still $\Theta(d)$, with a smaller implicit constant. Since the number of edges and average degree both change by at most a constant factor, this implies that the number of nodes also changes by only a constant factor, even as nodes are split in the second step. Thus there are $\Theta(n)$ nodes remaining once the process halts, and so G still satisfies the girth conjecture.

For the last property (counting cycles), let π be a path of length $2c$ with $e \in \pi$, and let v be any other node. The total number of ways to choose (π, v) is bounded by

$$\begin{aligned} & \underbrace{n}_{\text{choose } v} \cdot \underbrace{2c}_{\text{choose position of } e \text{ in } \pi} \cdot \underbrace{O\left(n^{1/k}\right)^{2c-1}}_{\text{extend path by an edge } 2c-1 \text{ times}} \\ & \leq O\left(cn^{1+(2c-1)/k}\right) \\ & = O\left(cn^{\frac{k+2c-1}{k}}\right). \end{aligned}$$

We then argue that there can be at most one $(2k + 2c)$ -cycle that contains π as a subpath and v as its antipodal node opposite π (with exactly k steps from either endpoint of π to v). Let s, t be the endpoint nodes of π , and notice that any such cycle can be viewed as π concatenated with a simple $s \rightsquigarrow v$ path of length k and a simple $t \rightsquigarrow v$ path of length k . However, there can be at most one simple $s \rightsquigarrow v$ (or $t \rightsquigarrow v$) path of length k , since if there are two such paths then they will imply a cycle of length $\leq 2k$, contradicting the girth of G . Thus there can be at most one cycle of length exactly $(2k + 2c)$ that contains π as a subpath

and v as its antipodal node opposite π , and so our counting of choices for (π, v) gives an upper bound for the number of cycles that contain e .

Finally, note that any particular $(2k + 2c)$ -cycle can be witnessed by exactly $2k + 2c$ different choices of (π, v) , since v can be any node on the cycle and π is then the path opposite v . So the total number of $(2k + 2c)$ -cycles is

$$O\left(\frac{c}{c+k} \cdot n^{\frac{k+2c-1}{k}}\right) \leq O\left(n^{\frac{k+2c-1}{k}}\right),$$

as claimed. \square

2.2 Lower Bound Construction

We will next describe our lower bound construction:

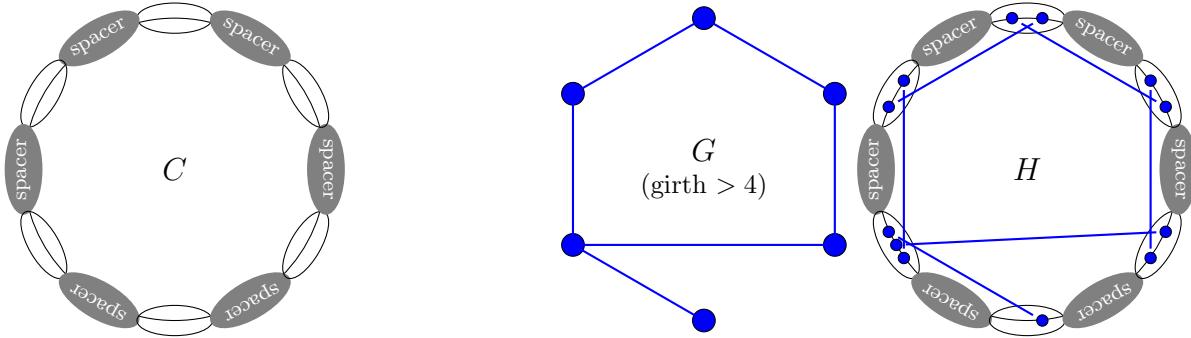


Figure 1: We construct our lower bound graph H by mapping the nodes of a girth conjecture graph (with parameter $k - 1$) G into the “clusters” of a large cycle C , which are separated by ample spacers, and then mapping the edges of G to random edges between clusters.

Lower Bound Construction.

- Let G be an n -node graph satisfying the girth conjecture with respect to parameter $k - 1$, and with the additional properties listed in Lemma 2.1 (bipartite, approximately regular, bounded number of $(2k + 2c)$ -cycles).
- Let C be a cycle on $N := 4k\varepsilon^{-1}n$ nodes in which all edges have weight 1. We will refer to C as the *spanning cycle*, or sometimes SC for brevity.
- Partition C into node-disjoint subpaths that alternately contain $k\varepsilon^{-1}$ nodes (called *clusters*) and $3k\varepsilon^{-1}$ nodes (called *spacers*). Note that there are n clusters and n spacers in total.
- Choose an arbitrary correspondence between the nodes of G and the clusters of C ; let us write X_v for the cluster assigned to a node v .
- For each edge $(u, v) \in E(G)$, place one edge of weight ε^{-1} between a uniform-random node pair in $X_u \times X_v$.
- For each cycle X in G of normalized weight $w^*(X) \leq (1 + \varepsilon)(2k)$, delete all non-SC edges in X from the graph.

We will let H be the final graph.

This construction clearly gives a graph H of weighted girth $> (1 + \varepsilon)(2k)$, since in the final step we explicitly remove edge(s) from every cycle of smaller normalized weight (and the spanning cycle itself has normalized weight $N \gg (1 + \varepsilon) \cdot 2k$). The lightness of H is given by:

LEMMA 2.2. Suppose that only a constant fraction of the non-SC edges are removed in the final step of the construction, due to participating in cycles of small normalized weight. Then the lightness of H is

$$\ell(H) = \Omega\left(\varepsilon^{1/(k-1)} \cdot \frac{N^{1/(k-1)}}{k}\right).$$

Proof. Since we assume that only a constant fraction of the weight is removed in the final step, it suffices to calculate the lightness of H just before this step is performed. This is:

$$\begin{aligned} \ell(H) &:= \frac{w(H)}{w(\text{mst}(H))} \\ &\geq \Omega\left(\frac{w(H \setminus C)}{w(C)}\right) \\ &= \Theta\left(\frac{\varepsilon^{-1}|E(G)|}{N}\right) \\ &= \Theta\left(\frac{\varepsilon^{-1}n^{1+1/(k-1)}}{N}\right) \\ &= \Theta\left(\frac{\varepsilon^{-1}\left(\frac{\varepsilon N}{k}\right)^{1+1/(k-1)}}{N}\right) \\ &= \Theta\left(\varepsilon^{1/(k-1)}N^{1/(k-1)}k^{-1}\right). \end{aligned}$$

□

Note that this lightness bound is *increasing* with ε , and so currently we would like to choose ε as large as possible, since this improves both the weighted girth and the lightness of our construction. The catch is that we need to satisfy the hypothesis that only a constant fraction of the non-SC edges are removed in the final step, and we will see in the following analysis that this requires us to enforce an upper bound on ε . Thus, our final setting of ε will be right at this upper bound.

2.3 Cycle Analysis In order to reason about the number of edges that are removed in the final step of our construction of H , we need some technical lemmas that inspect the structure of the cycles just before this step. Let H' be the graph just before the last step of the construction is performed (so $H \subseteq H'$).

LEMMA 2.3. Assuming $\varepsilon < \frac{1}{2}$, every cycle X in H' of normalized weight $w^*(X) \leq 2k(1 + \varepsilon)$ contains at least $2k$ and at most $2k(1 + \varepsilon)$ non-SC edges.

Proof. We show the contrapositive. If X has $> 2k(1 + \varepsilon)$ non-SC edges, then its normalized weight is

$$\begin{aligned} w^*(X) &:= \frac{w(X)}{\max_{e \in X} w(e)} \\ &> \frac{2k(1 + \varepsilon) \cdot \varepsilon^{-1}}{\varepsilon^{-1}} \\ &= 2k(1 + \varepsilon). \end{aligned}$$

On the other hand, suppose that X has $< 2k$ non-SC edges. Note that G has girth $\geq 2k$, since (by assumption) it has no cycles of length $\leq 2k - 2$, and (since it is bipartite) it also has no cycles of length $2k - 1$. Thus the non-SC edges of X do *not* correspond to a cycle in G . That means there are two adjacent non-SC edges $e_1, e_2 \in X$ that do not share a cluster, and so X must include a path of SC edges between them that goes between their respective clusters. This path includes at least one spacer, and due to this

spacer we have

$$\begin{aligned} w^*(X) &:= \frac{w(X)}{\max_{e \in X} w(e)} \\ &\geq \frac{3k\varepsilon^{-1}}{\varepsilon^{-1}} \\ &> 2k(1 + \varepsilon). \end{aligned}$$

since we assume $\varepsilon < \frac{1}{2}$.

□

In the following, for any cycle X in G , we will say that the *corresponding* cycle in H' is the one that contains the image of each edge in X , as well as the shortest SC path between any two adjacent edge-images joining their endpoints.

LEMMA 2.4. *Let $c \geq 0$ be an integer and let X be a $(2k + 2c)$ -cycle in G . Then the probability that the corresponding cycle X' in H' has normalized weight $\leq (1 + \varepsilon)2k$ is at most*

$$\frac{\Theta(\varepsilon)^{2k+2c}}{(2k + 2c)!}.$$

Proof. Let σ be the total number of SC edges used by X' . The normalized weight of X' is

$$\begin{aligned} w^*(X') &:= \frac{w(X')}{\max_{e \in X'} w(e)} \\ &= \frac{(2k + 2c)\varepsilon^{-1} + \sigma}{\varepsilon^{-1}} \\ &= 2k + 2c + \varepsilon\sigma. \end{aligned}$$

This is $\leq (1 + \varepsilon)2k$ iff $\sigma \leq 2k - 2c\varepsilon^{-1}$. So our goal is to bound the probability that the random mapping of the edges of X' into C yields $2k - 2c\varepsilon^{-1}$ edges in total, along the SC paths joining their endpoints.

As a first step, let us count the number of ways to partition exactly $2k - 2c\varepsilon^{-1}$ steps among $2k + 2c$ clusters, plus a “remainder” bucket (to account for the possibility that fewer steps are actually used within clusters). This is a standard application of the stars-and-bars counting formula, which gives

$$\binom{2k - 2c\varepsilon^{-1} + 2k + 2c}{2k + 2c} \leq \frac{(4k)^{2k+2c}}{(2k + 2c)!}$$

ways. For each such partition, the probability that it is realized by the mapping of X into H' can be bounded by treating the first endpoint mapped to each cluster as fixed, and then acknowledging that there are at most two nodes to which we could map the second endpoint in that cluster that will cause exactly the selected number of in-cluster steps to be used. So the probability is at most

$$\left(\frac{2}{k\varepsilon^{-1}}\right)^{2k+2c}.$$

Multiplying the previous two terms, we get a bound of

$$\frac{\Theta(\varepsilon)^{2k+2c}}{(2k + 2c)!}.$$

□

LEMMA 2.5. *For any integer $0 \leq c \leq k$, if*

$$\varepsilon \leq O\left(k \cdot n^{-\frac{k+2c}{(k-1)(2k+2c)}}\right)$$

with a small enough implicit constant, then for each non-SC edge e , the expected number of cycles X' in H' that have $e \in X'$, exactly $2k + 2c$ non-SC edges, and normalized weight $w^(X') \leq (1 + \varepsilon)2k$ is at most $\frac{1}{4^{c+1}}$.*

Proof. First, let us bound the number of $(2k + 2c)$ -cycles in G . Recall that G is a girth conjecture graph with parameter $k - 1$. So, instead viewing these as $(2(k - 1) + 2(c + 1))$ -cycles and then applying Lemma 2.1, the number of $(2k + 2c)$ -cycles in G that contain e is bounded by

$$O\left(n^{\frac{k+2c}{k-1}}\right).$$

By Lemma 2.4, the probability that each of these cycles in G is mapped to a cycle of small normalized weight in H' is at most

$$\frac{\Theta(\varepsilon)^{2k+2c}}{(2k + 2c)!}.$$

We would like to set ε sufficiently small that the expected number of such cycles is at most $1/4^{c+1}$; that is,

$$O\left(n^{\frac{k+2c}{k-1}}\right) \cdot \frac{\Theta(\varepsilon)^{2k+2c}}{(2k + 2c)!} \leq \frac{1}{4^{c+1}}.$$

Solving for ε , we get

$$\begin{aligned} n^{\frac{k+2c}{k-1}} \cdot \Theta(\varepsilon)^{2k+2c} &\leq O((2k + 2c)!) \\ \Theta(\varepsilon)^{2k+2c} &\leq O\left((2k + 2c)! \cdot n^{-\frac{k+2c}{k-1}}\right) \\ \varepsilon &\leq O\left((2k + 2c)!^{\frac{1}{2k+2c}} \cdot n^{-\frac{k+2c}{(k-1)(2k+2c)}}\right) \\ \varepsilon &\leq O\left(k \cdot n^{-\frac{k+2c}{(k-1)(2k+2c)}}\right). \end{aligned}$$

□

2.4 Proof Wrapup The previous lemma dictates our final choice of ε , but note that it expresses its bound as a function of n (the number of nodes in G), rather than N (the number of nodes in H). To rearrange the bound as a function of N , we calculate:

$$\begin{aligned} \varepsilon &= \Theta\left(k \cdot n^{-\frac{k+2c}{(k-1)(2k+2c)}}\right) \\ \varepsilon &= \Theta\left(k \cdot \left(\frac{\varepsilon N}{k}\right)^{-\frac{k+2c}{(k-1)(2k+2c)}}\right) \\ \varepsilon^{1+\frac{k+2c}{(k-1)(2k+2c)}} &= \Theta\left(k \cdot N^{-\frac{k+2c}{(k-1)(2k+2c)}} \cdot k^{\Theta(1/k)}\right) \\ \varepsilon^{\frac{2k^2+2kc-k}{(k-1)(2k+2c)}} &= \Theta\left(k \cdot N^{-\frac{k+2c}{(k-1)(2k+2c)}}\right) \\ \varepsilon^{2k^2+2kc-k} &= \Theta\left(k^{(k-1)(2k+2c)} \cdot N^{-(k+2c)}\right) \\ \varepsilon &= \Theta\left(k^{\frac{(k-1)(2k+2c)}{2k^2+2kc-k}} \cdot N^{-\frac{(k+2c)}{2k^2+2kc-k}}\right). \end{aligned}$$

To wrap up the proof, we need to look more carefully at the setting of c . Notice that Lemma 2.3 gives an upper bound on the values of c that we need to consider: that is, so long as $c > k\varepsilon$, there are *no* cycles that use $2k + 2c$ non-SC edges and which have weighted girth $\leq (1 + \varepsilon) \cdot 2k$. So the appropriate setting of ε is given by the above formula, with $c := \lfloor k\varepsilon \rfloor$.

Under this setting, for any edge e , by Lemma 2.5 the expected number of cycles in H' that contain e and have normalized weight $\leq (1 + \varepsilon) \cdot 2k$ is bounded by

$$\sum_{i=0}^{\lfloor k\varepsilon \rfloor} \frac{1}{4^{i+1}} \leq \frac{1}{2}.$$

By Markov's inequality, this implies that e survives in the final step of our construction with constant probability, and so the conditions of Lemma 2.2 are satisfied and the analysis is complete.

Finally, let us acknowledge that our original main result was phrased more restrictively (constant k) but also had a simpler-looking setting of ε . For fixed k , our setting of ε is $\ll 1/k$, which means that the appropriate setting of ε is the one that takes $c = 0$. With this restriction in place, we can continue to simplify our setting of ε :

$$\begin{aligned}\varepsilon &= \Theta\left(k^{\frac{(k-1)(2k)}{2k^2-k}} \cdot N^{-\frac{k}{2k^2-k}}\right) \\ \varepsilon &= \Theta\left(k \cdot N^{-\frac{1}{2k-1}}\right),\end{aligned}$$

which is the setting in our main result discussed in the introduction. To then calculate our lower bound, by Lemma 2.2 we have

$$\begin{aligned}\ell(H) &= \Omega_k\left(\varepsilon^{\frac{1}{k-1}} \cdot N^{\frac{1}{k-1}}\right) \\ &= \Omega_k\left(\left(N^{-\frac{1}{2k-1}}\right)^{\frac{1}{k-1}} \cdot N^{\frac{1}{k-1}}\right) \\ &= \Omega_k\left(N^{\frac{1}{k}} \cdot N^{\frac{1}{k-1} - \frac{1}{k} - \frac{1}{(2k-1)(k-1)}}\right) \\ &= \Omega_k\left(N^{\frac{1}{k}} \cdot N^{\frac{1}{k(2k-1)}}\right).\end{aligned}$$

This bound is the one claimed in Corollary 1.1. Finally, since $\varepsilon = \Theta_k(N^{-\frac{1}{2k-1}})$, we can rewrite this bound as

$$\Omega_k\left(N^{\frac{1}{k}} \cdot \varepsilon^{-\frac{1}{k}}\right),$$

completing the proof.

3 Proof of Technical Lower Bound

We now prove Theorem 1.5, establishing limitations to improving our lower bound further. For this proof, let $\gamma(n, 2k)$ denote the maximum possible number of edges in an n -node graph of girth $> 2k$. So the girth conjecture with parameter k states that

$$\gamma(n, 2k) \geq \Omega(n^{1+1/k}).$$

Prior work on light spanners has established upper bounds that are *independent* of the girth conjecture, in the sense that their lightness is a function of γ . Specifically:

THEOREM 3.1. ([9, 12]) *For all positive integers n, k and all $\varepsilon > 0$, every n -node graph has a $(1+\varepsilon)(2k-1)$ spanner H of lightness*

$$\ell(H \mid G) \leq O\left(\varepsilon^{-1} \cdot \frac{\gamma(n, 2k)}{n}\right).$$

We will use this in our proof of Theorem 1.5.

Proof. [Proof of Theorem 1.5] Our proof of the girth conjecture is by induction on k ; we may incur constant factors at each level of the induction, since its depth is limited to k and we treat k as a constant. The base case is $k = 2$, where the girth conjecture is already proved [16].

In the inductive step, since Theorem 1.3 is only conditional on the girth conjecture with parameter $k-1$, by the inductive hypothesis we may apply it in our construction to get an (unconditional) lower bound on lightness of

$$\ell(H) \geq \Omega\left(\varepsilon^{-1/k} n^{1/k}\right).$$

Comparing this lower bound to the upper bound in Theorem 3.1, we must have

$$\Omega\left(\varepsilon^{-1/k} n^{1/k}\right) \leq O\left(\varepsilon^{-1} \cdot \frac{\gamma(n, 2k)}{n}\right).$$

Now, if we assume that Theorem 1.3 extends to a parameter regime where ε is a constant (perhaps depending on k), then we may set ε to be a constant and rearrange to get

$$\Omega\left(n^{1+1/k}\right) \leq \gamma(n, 2k),$$

implying the girth conjecture for parameter k . Alternately, if we can improve the left-hand side dependence from $\varepsilon^{-1/k}$ to ε^{-1} , then the ε terms cancel and rearranging in the same way again implies the girth conjecture for parameter k , completing the inductive step. \square

References

- [1] Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020.
- [2] Stephen Alstrup, Søren Dahlgaard, Arnold Filtser, Morten Stöckel, and Christian Wulff-Nilsen. Constructing light spanners deterministically in near-linear time. *Theoretical Computer Science*, 907:82–112, 2022.
- [3] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.
- [4] Greg Bodwin. An alternate proof of near-optimal light spanners. In *2024 Symposium on Simplicity in Algorithms (SOSA)*, pages 39–55. SIAM, 2024.
- [5] Glencora Borradaile, Hung Le, and Christian Wulff-Nilsen. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 767–778. IEEE, 2017.
- [6] Glencora Borradaile, Hung Le, and Christian Wulff-Nilsen. Greedy spanners are optimal in doubling metrics. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2371–2379. SIAM, 2019.
- [7] Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. New sparseness results on graph spanners. In *Proceedings of the eighth annual symposium on Computational geometry*, pages 192–201. ACM, 1992.
- [8] Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. *ACM Transactions on Algorithms (TALG)*, 14(3):1–15, 2018.
- [9] Michael Elkin, Ofer Neiman, and Shay Solomon. Light spanners. *SIAM Journal on Discrete Mathematics*, 29(3):1312–1321, 2015.
- [10] Paul Erdős. Extremal problems in graph theory. In *Proceedings of the Symposium on Theory of Graphs and its Applications*, page 2936, 1963.
- [11] Arnold Filtser and Shay Solomon. The greedy spanner is existentially optimal. *SIAM Journal on Computing*, 49(2):429–447, 2020.
- [12] Hung Le and Shay Solomon. A unified framework for light spanners. In *Proceedings of the 55th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 2023.
- [13] David Peleg and Jeffrey Ullman. An optimal synchronizer for the hypercube. *SIAM Journal on Computing (SICOMP)*, 18(4):740—747, 1989.
- [14] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)*, 36(3):510–530, 1989.
- [15] Jacques Tits. Sur la trialité et certains groupes qui s’ en déduisent. *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, 2(1):14–60, 1959.
- [16] R. Wenger. Extremal graphs with no c4’s, c6’s, or c10’s. *J. Combin. Theory Ser. B*, 52(1):113–116, 1991.